CLAIMS

What is claimed is:

1. A method of discovering nodes, comprising:

probing an $M \times N$ torus interconnection fabric, wherein M and N are integer values and said interconnection fabric includes a first plurality of nodes forming an x-axis and a second plurality of nodes forming a y-axis; and

identifying a location of a first node relative to the x and y axes.

2. The method of Claim 1, wherein said probing the $M \times N$ torus interconnection fabric further comprises:

determining the number of nodes in the x-direction of the interconnection fabric; and

determining the number of nodes in the y-direction of the interconnection fabric.

3. The method of Claim 1, further comprising:

after identifying the location of the first node relative to the x and y axes, assigning an identification code to each of the nodes in the interconnection fabric.

4. The method of Claim 1, wherein:

at least one of the nodes in the interconnection fabric is connected to a storage device.

5. The method of Claim 1, wherein:

the probing the interconnection fabric comprises sending probe messages from the first node to query a first set of nodes in the interconnection fabric to identify at least one of the first plurality of nodes forming the x-axis and at least one of the second plurality of nodes forming the y-axis.

6. The method of Claim 5, wherein:

each of the nodes in the interconnection fabric has an associated origin code; and

said identifying the location of the first node relative to the *x* and *y* axes comprises, in response to sending the probe messages from the first node to query the first set of nodes, receiving the origin codes from the first set of nodes.

7. The method of Claim 6, wherein:

said identifying the location of the first node relative to the *x* and *y* axes comprises storing the received origin codes from the first set of nodes in a first origin code mapping.

8. The method of Claim 6, wherein each of said first plurality of nodes forming the *x*-axis has the associated origin code set to a first value, each of said second plurality of nodes forming the *y*-axis has the associated origin code set to the first value, and each of the nodes not forming either the *x* or *y* axes has the associated origin code set to a second value.

9. The method of Claim 8, wherein each of the nodes in the interconnection fabric includes a one-bit memory for storing the associated origin code.

10. The method of Claim 8, wherein a node at an intersection of the *x* and *y* axes forms an origin node and has the associated origin code set to the second value.

11. The method of Claim 10, wherein:

said sending probe messages from the first node to query the first set of nodes in the interconnection fabric comprises:

sending probe messages from the first node to query nodes in a row of the interconnection fabric; and

sending probe messages from the first node to query nodes in a column of the interconnection fabric; and

said identifying the location of the first node relative to the *x* and *y* axes comprises:

receiving the origin codes from the nodes in the row of the interconnection fabric to identify a change of origin code indicating the location of the *y* axis; and

receiving the origin codes from the nodes in the column of the interconnection fabric to identify a change of origin code indicating the location of the *x* axis.

12. The method of Claim 6, further comprising:

probing of the interconnection fabric by sending probe messages from a second node to query a second set of nodes in the interconnection fabric to identify at least one of the first plurality of nodes forming the *x*-axis and at least one of the second plurality of nodes forming the *y*-axis; and

identifying a location of the second node relative to the *x* and *y* axes by receiving the origin codes from the second set of nodes.

13. The method of Claim 12, further comprising:

comparing the origin codes received from the first set of nodes to the origin codes received from the second set of nodes.

14. The method of Claim 12, further comprising:

using the origin codes received by the node having a higher priority level to assign an identification code to each of the nodes in the interconnection fabric.

15. A computer system, comprising:

an $M \times N$ array of nodes, wherein M and N are integer values; and

a plurality of interconnects connecting the $M \times N$ array;

wherein:

a first plurality of nodes in the $M \times N$ array form an x-axis in the $M \times N$ array;

a second plurality of nodes in the $M \times N$ array form a y-axis in the $M \times N$ array; and

a first node in the $M \times N$ array is configured to probe the $M \times N$ array to identify a location of the first node relative to the x-axis and the y-axis.

16. The system of Claim 15, wherein:

the first node is further configured to determine the number of nodes in the x-direction of the $M \times N$ array and to determine the number of nodes in the y-direction of the $M \times N$ array.

17. The system of Claim 15, wherein:

the first node is further configured to, after probing the $M \times N$ array to identify the location of the first node relative to the x and y axes, assign an identification code to each of the nodes in the $M \times N$ array.

18. The system of Claim 15, wherein:

at least one of the nodes in the $M \times N$ array is connected to a storage device.

19. The system of Claim 15, wherein:

the first node is further configured to probe the $M \times N$ array by sending probe messages to query a first set of nodes in the $M \times N$ array to identify at least one of the first plurality of nodes forming the x-axis and at least one of the second plurality of nodes forming the y-axis.

20. The system of Claim 19, wherein:

each of the nodes in the $M \times N$ array includes an associated origin code; and

the first node is further configured to receive the origin codes from the first set of nodes in response to sending probe messages to query the first set of nodes.

21. The system of Claim 20, wherein:

the first node is further configured to store the received origin codes from the first set of nodes in a first mapping.

22. The system of Claim 20, wherein:

each of the first plurality of nodes forming the x-axis have the origin code set to a first value;

each of the second plurality of nodes forming the y-axis have the origin code set to the first value; and

each of the nodes in the $M \times N$ array not defined as forming either the x or y axes has the origin code set to a second value.

23. The system of Claim 22, wherein:

each of the nodes in the $M \times N$ array include a one-bit memory to store the origin code.

24. The system of Claim 22, further comprising:

a node at an intersection of the x and y axes is defined as an origin node by setting the origin code for the origin node to the second value.

25. The system of Claim 24, wherein:

said first node is further configured to send probe messages to query the first set of nodes in the interconnection fabric by sending probe messages to query nodes in a row of the interconnection fabric, and sending probe messages to query nodes in a column of the interconnection fabric; and

said first node is further configured to identify the location of the first node relative to the x and y axes by receiving the origin codes from the nodes in the row of the interconnection fabric to identify a change of origin code indicating the location of the y axis, and receiving the origin codes from the nodes in the column of the interconnection fabric to identify a change of origin code indicating the location of the x axis.

26. The system of Claim 19, further comprising:

a second node in the $M \times N$ array configured to probe the $M \times N$ array to identify a location of the second node relative to the x and y axes by sending probe messages to query a second set of nodes in the $M \times N$ array to identify at least some of the first plurality of nodes forming the x-axis and at least some of the second plurality of nodes forming the y-axis, and receiving the origin codes from the second set of nodes.

27. The system of Claim 26, wherein:

the first node is further configured to transmit to the second node the origin codes received from the first set of nodes;

the second is further configured to transmit to the first node the origin codes received from the second set of nodes; and

the first and second nodes are configured to compare the origin codes received from the first set of nodes to the origin codes received from the second set of nodes.

28. The system of Claim 26, wherein:

each of the first and second nodes have a priority level, and the node having the higher priority level uses the origin codes received by the node having the higher priority level to assign an identification code to each of the nodes in the $M \times N$ array.

29. The system of Claim 15, wherein the first node is a CPU node.

29. An article of manufacture including code for discovering nodes, wherein the code causes operations to be performed comprising:

probing an $M \times N$ torus interconnection fabric, wherein M and N are integer values and said interconnection fabric includes a first plurality of nodes forming an x-axis and a second plurality of nodes forming a y-axis; and

identifying a location of a first node relative to the x and y axes.

30. The article of manufacture of Claim 29, further comprising:

determining the number of nodes in the x-direction of the interconnection fabric; and

determining the number of nodes in the y-direction of the interconnection fabric.

31. The article of manufacture of Claim 29, further comprising:

after identifying the location of the first node relative to the x and y axes, assigning an identification code to each of the nodes in the interconnection fabric.

32. The article of manufacture of Claim 29, wherein:

at least one of the nodes in the interconnection fabric is connected to a storage device.

33. The article of manufacture of Claim 29, wherein:

the probing the interconnection fabric comprises sending probe messages from the first node to query a first set of nodes in the interconnection fabric to identify at least one of the first plurality of nodes forming the *x*-axis and at least one of the second plurality of nodes forming the *y*-axis.

34. The article of manufacture of Claim 33, wherein:

each of the nodes in the interconnection fabric has an associated origin code; and

said identifying the location of the first node relative to the *x* and *y* axes comprises, in response to sending the probe messages from the first node to query the first set of nodes, receiving the origin codes from the first set of nodes.

35. The article of manufacture of Claim 34, wherein:

said identifying the location of the first node relative to the *x* and *y* axes comprises storing the received origin codes from the first set of nodes in a first origin code mapping.

36. The article of manufacture of Claim 34, wherein each of said first plurality of nodes forming the *x*-axis has the associated origin code set to a first value, each of said second plurality of nodes forming the *y*-axis has the associated origin code set to the first value, and each of the nodes not forming either the *x* or *y* axes has the associated origin code set to a second value.

37. The article of manufacture of Claim 36, wherein:

a node at an intersection of the *x* and *y* axes forms an origin node and has the associated origin code set to the second value;

said sending probe messages from the first node to query the first set of nodes in the interconnection fabric comprises:

sending probe messages from the first node to query nodes in a row of the interconnection fabric; and

sending probe messages from the first node to query nodes in a column of the interconnection fabric; and

said identifying the location of the first node relative to the *x* and *y* axes comprises:

receiving the origin codes from the nodes in the row of the interconnection fabric to identify a change of origin code indicating the location of the *y* axis; and

receiving the origin codes from the nodes in the column of the interconnection fabric to identify a change of origin code indicating the location of the *x* axis.

38. The article of manufacture of Claim 34, further comprising:

probing of the interconnection fabric by sending probe messages from a second node to query a second set of nodes in the interconnection fabric to identify at least one of the first plurality of nodes forming the *x*-axis and at least one of the second plurality of nodes forming the *y*-axis; and

identifying a location of the second node relative to the *x* and *y* axes by receiving the origin codes from the second set of nodes.

39. The article of manufacture of Claim 38, further comprising:

comparing the origin codes received from the first set of nodes to the origin codes received from the second set of nodes.

40. The article of manufacture of Claim 39, further comprising:

using the origin codes received by the node having a higher priority level to assign an identification code to each of the nodes in the interconnection fabric.

41. A method of discovering nodes, comprising:

probing an $M \times N$ torus interconnection fabric, wherein M and N are integer values and said probing comprises sending probe messages from a first node to query a first set of nodes in the interconnection fabric to identify at least one of a first plurality of nodes forming an x-axis and at least one of a second plurality of nodes forming a y-axis;

identifying a location of the first node relative to the x and y axes, wherein each of the nodes in the interconnection fabric has an associated origin code and said identifying comprises, in response to sending the probe messages from the first node to query the first set of nodes, receiving the origin codes from the first set of nodes;

generating an observed mapping of the nodes in the interconnection fabric showing a location of a first node relative to an x-axis of the fabric and relative to a y-axis of the fabric based on the origin codes received from the first set of nodes;

comparing the observed mapping of the nodes to a set of expected mappings; and

identifying the expected mapping which is most similar to the observed mapping.

42. The method of Claim 41, further comprising:

assigning identification codes to each of the nodes based on the identified mapping which is most similar to the observed mapping.

43. The method of Claim 41, wherein:

the set of expected mappings contains a set of all possible valid mappings of origin codes for the $M \times N$ fabric; and

said comparing the observed mapping of the nodes to the set of expected mappings comprises comparing the received origin codes in the observed mapping to the origin codes in each expected mapping in the set of expected mappings.